ABSTRACT

A parallel processing network permits processes to be spawned based on the availability of various network features. Such features may include the type of CPU's in the network, the number of CPU's per machine and other network resources. A user can select either a process group file method of process spawning or an automatic spawning method. In the automatic method, the user specifies various criteria related to how the processes are to be spawned such as the desired number of processes to be spawned, the type of CPUs to which the processes are to be spawned, the maximum number of processes to be started on any one machine and other information as desired. The spawning routine preferably runs on a root machine and accesses a process scheduler which provides the current network configuration. If CPUs and machines are available (i.e., operational) that match the user's criteria as determined by access to a process scheduler, the user desired number of processes is spawned to the CPUs and machines that match the criteria. If there are not enough CPUs and/or machines that match the user's criteria, the spawning routine decreases the number of processes from the user desired number of processes, and spawns processes to as many CPUs and machines that otherwise match the user's criteria. As such, the parallel processing network advantageously permits processes to be spawned automatically without requiring the user to have a detailed understanding of the available network features.